arXiv:1702.05462v1 [stat.ME] 17 Feb 2017

Objective Bayesian Analysis for Change Point Problems

Laurentiu Hinoveanu*[1], Fabrizio Leisen†[1], and Cristiano Villa‡[1]

[1]School of Mathematics, Statistics and Actuarial Science, University of Kent

Abstract

In this paper we present an objective approach to change point analysis. In particular, we look at the problem from two perspectives. The first focuses on the definition of an objective prior when the number of change points is known a priori. The second contribution aims to estimate the number of change points by using an objective approach, recently introduced in the literature, based on losses. The latter considers change point estimation as a model selection exercise. We show the performance of the proposed approach on simulated data and on real data sets.

Keywords: Change point; Discrete parameter space; Loss-based prior; Model selection; Objective Bayes.

1 Introduction

There are several practical scenarios where it is inappropriate to assume that the distribution of the observations does not change. For example, financial datasets can exhibit alternate behaviours due to crisis periods. In this case it is sensible to assume changes in the underlying distribution. The change in the distribution can be either in the value of one or more of the parameters

*lch36@kent.ac.uk
†F.Leisen@kent.ac.uk
‡C.Villa-88@kent.ac.uk

or, more in general, on the family of the distribution. In the latter case, one may deem appropriate to consider a normal density for the stagnation periods, while a Student t, with relatively heavy tails, may be more suitable to represents observations in the more turbulent data of a crisis. The task of identifying if, and when, one or more changes have occurred is not trivial and requires appropriate methods to avoid detection of a large number of changes or, at the opposite extreme, seeing no changes at all. Whilst the literature covering change point analysis from a Bayesian perspective is vast when prior distributions are elicited, the documentation referring to analysis under minimal prior information (i.e. objective Bayes), is very limited. In fact, to the best of our knowledge, only two papers tackled the problem from an objective point of view: Moreno et al. (2005) and Girón et al. (2007). The former discusses the single change point problem in a model selection setting, whilst the latter, which is an extension of the former, tackles the multivariate change point problem in the context of linear regression models.

This work aims to contribute to the methodology for change point analysis under the assumption that the information about the number of change points and their location is minimal. First, we discuss the definition of an objective prior for change point location, both for single and multiple changes, assuming the number of changes is known a priori. Then, we define a prior on the number of change points via a model selection approach. Here, we assume that the change point coincides with one of the observations. As such, given $X_1, X_2, \ldots, X_n$ data points, the change point location is discrete. To the best of our knowledge, the sole objective approach to define prior distributions on discrete spaces is the one introduced by Villa and Walker (2015a).

To illustrate their idea, consider a probability distribution $f(x|m)$, where $m \in \mathbb{M}$ is a discrete parameter. Then, the prior $\pi(m)$ is obtained by objectively measuring what is lost if the value m is removed from the parameter space, and it is the true value. According to Berk (1966), if a model is misspecified, the posterior distribution asymptomatically accumulates on the model which is the most similar to the true one, where the similarity is measured in terms of the Kullback–Leibler (KL) divergence. Therefore, $D_{KL}(f(\cdot|m)\|f(\cdot|m'))$, where m' is the parameter characterising the nearest model to $f(x|m)$, represents the utility of keeping m. The objective prior is then obtained by linking the aforementioned utility via the self-information loss:

$$\pi(m) \propto \exp\left\{\min_{m' \neq m} D_{KL}(f(\cdot|m)\|f(\cdot|m'))\right\} - 1, \tag{1}$$

where the Kullback–Leibler divergence (Kullback and Leibler, 1951) from the sampling distribution with density $f(x|m)$ to the one with density $f(x|m')$ is defined as:

$$D_{KL}(f(\cdot|m)\|f(\cdot|m')) = \int_{\mathcal{X}} f(x|m) \cdot \log\left[\frac{f(x|m)}{f(x|m')}\right] \mathrm{d}x.$$

Throughout the paper, the objective prior defined in equation (1) will be referenced as the loss-based prior. This approach is used to define an objective prior distribution when the number of change points is known a priori. To obtain a prior distribution for the number of change points, we adopt a model selection approach based on the results in Villa and Walker (2015*b*), where a method to define an objective prior on the space of models is proposed. To illustrate, let us consider k Bayesian models:

$$M_j = \{f_j(x|\theta_j), \pi_j(\theta_j)\} \qquad j \in \{1, 2, \ldots, k\}, \tag{2}$$

where $f_j(x|\theta_j)$ is the sampling density characterised by θ_j and $\pi_j(\theta_j)$ represents the prior on the model parameter.

Assuming the prior on the model parameter, $\pi_j(\theta_j)$, is proper, the model prior probability $\Pr(M_j)$ is proportional to the expected minimum Kullback–Leibler divergence from M_j, where the expectation is considered with respect to $\pi_j(\theta_j)$. That is:

$$\Pr(M_j) \propto \exp\left\{\mathbb{E}_{\pi_j}\left[\inf_{\theta_i, i\neq j} D_{KL}(f_j(x|\theta_j)\|f_i(x|\theta_i))\right]\right\} \qquad j = 1, \ldots, k. \tag{3}$$

The model prior probabilities defined in equation (3) can be employed to derive the model posterior probabilities through:

$$\Pr(M_i|x) = \left[\sum_{j=1}^{k} \frac{\Pr(M_j)}{\Pr(M_i)} B_{ji}\right]^{-1}, \tag{4}$$

where B_{ji} is the Bayes factor between model M_j and model M_i, defined as

$$B_{ji} = \frac{\int f_j(x|\theta_j)\pi_j(\theta_j)\,\mathrm{d}\theta_j}{\int f_i(x|\theta_i)\pi_i(\theta_i)\,\mathrm{d}\theta_i},$$

with $i \neq j \in \{1, 2, \ldots, k\}$.

The paper is structured as follows. In Section 2 we establish the way we set objective priors on both single and multiple change point locations. Section

3 shows how we define the model prior probabilities for the number of change point locations. Illustrations of the model selection exercise are provided in Sections 4 and 5, where we work with simulated and real data, respectively. Section 6 is dedicated to final remarks.

2 Objective Prior on the Change Point Locations

This section is devoted to the derivation of the loss-based prior when the number of change points is known a priori. Specifically, let k be the number of change points and $m_1 < m_2 < \ldots < m_k$ their locations. We introduce the idea in the simple case where we assume that there is only one change point in the dataset (see Section 2.1). Then, we extend the results to the more general case where multiple change points are assumed (see Section 2.2).

2.1 Single Change Point

Here, we show that the loss-based prior for the single change point case coincides with the discrete uniform over the set $\{1, 2, \ldots, n-1\}$.

Let $\mathbf{X}^{(n)} = (X_1, \ldots, X_n)$ denote an n-dimensional vector of random variables, representing the random sample, and m be our single change point location, that is $m \in \{1, 2, \ldots, n-1\}$, such that

$$\begin{aligned} X_1, \ldots, X_m | \tilde{\theta}_1 &\overset{\text{i.i.d.}}{\sim} f_1(\cdot|\tilde{\theta}_1) \\ X_{m+1}, \ldots, X_n | \tilde{\theta}_2 &\overset{\text{i.i.d.}}{\sim} f_2(\cdot|\tilde{\theta}_2). \end{aligned} \tag{5}$$

Note that we assume that there is a change point in the series, as such the space of m does not include the case $m = n$. In addition, we assume that $\tilde{\theta}_1 \neq \tilde{\theta}_2$ when $f_1 = f_2$. The sampling density for the vector of observations $\mathbf{x}^{(n)} = (x_1, \ldots, x_n)$ is:

$$f(\mathbf{x}^{(n)}|m, \tilde{\theta}_1, \tilde{\theta}_2) = \prod_{i=1}^{m} f_1(x_i|\tilde{\theta}_1) \prod_{i=m+1}^{n} f_2(x_i|\tilde{\theta}_2). \tag{6}$$

Let $m' \neq m$. Then, the Kullback–Leibler divergence between the model

parametrised by m and the one parametrised by m' is:

$$D_{KL}(f(\mathbf{x}^{(n)}|m,\tilde{\theta}_1,\tilde{\theta}_2)\|f(\mathbf{x}^{(n)}|m',\tilde{\theta}_1,\tilde{\theta}_2)) = \int f(\mathbf{x}^{(n)}|m,\tilde{\theta}_1,\tilde{\theta}_2) \log\left(\frac{f(\mathbf{x}^{(n)}|m,\tilde{\theta}_1,\tilde{\theta}_2)}{f(\mathbf{x}^{(n)}|m',\tilde{\theta}_1,\tilde{\theta}_2)}\right) \mathrm{d}\mathbf{x}^{(n)}. \tag{7}$$

Without loss of generality, consider $m < m'$. In this case, note that

$$\frac{f(\mathbf{x}^{(n)}|m,\tilde{\theta}_1,\tilde{\theta}_2)}{f(\mathbf{x}^{(n)}|m',\tilde{\theta}_1,\tilde{\theta}_2)} = \prod_{i=m+1}^{m'} \frac{f_2(x_i|\tilde{\theta}_2)}{f_1(x_i|\tilde{\theta}_1)},$$

leading to

$$D_{KL}(f(\mathbf{x}^{(n)}|m,\tilde{\theta}_1,\tilde{\theta}_2)\|f(\mathbf{x}^{(n)}|m',\tilde{\theta}_1,\tilde{\theta}_2)) = \sum_{i=m+1}^{m'} \int f_2(x_i|\tilde{\theta}_2) \log\left(\frac{f_2(x_i|\tilde{\theta}_2)}{f_1(x_i|\tilde{\theta}_1)}\right) \mathrm{d}x_i. \tag{8}$$

On the right hand side of equation (8), we can recognise the Kullback–Leibler divergence from density f_2 to density f_1, thus getting:

$$D_{KL}(f(\mathbf{x}^{(n)}|m,\tilde{\theta}_1,\tilde{\theta}_2)||f(\mathbf{x}^{(n)}|m',\tilde{\theta}_1,\tilde{\theta}_2)) = (m'-m)D_{KL}(f_2(\cdot|\tilde{\theta}_2)\|f_1(\cdot|\tilde{\theta}_1)). \tag{9}$$

In a similar fashion, when $m > m'$, we have that:

$$D_{KL}(f(\mathbf{x}^{(n)}|m,\tilde{\theta}_1,\tilde{\theta}_2)\|f(\mathbf{x}^{(n)}|m',\tilde{\theta}_1,\tilde{\theta}_2)) = (m-m')D_{KL}(f_1(\cdot|\tilde{\theta}_1)\|f_2(\cdot|\tilde{\theta}_2)). \tag{10}$$

In this single change point scenario, we can consider m' as a perturbation of the change point location m, that is $m' = m \pm l$ where $l \in \mathbb{N}^*$, such that $1 \leq m' < n$. Then, taking into account equations (9) and (10), the Kullback–Leibler divergence becomes:

$$D_{KL}(f(\mathbf{x}^{(n)}|m,\tilde{\theta}_1,\tilde{\theta}_2)\|f(\mathbf{x}^{(n)}|m',\tilde{\theta}_1,\tilde{\theta}_2)) = \begin{cases} l \cdot D_{KL}(f_2(\cdot|\tilde{\theta}_2)\|f_1(\cdot|\tilde{\theta}_1)), & \text{if } m < m' \\ l \cdot D_{KL}(f_1(\cdot|\tilde{\theta}_1)\|f_2(\cdot|\tilde{\theta}_2)), & \text{if } m > m', \end{cases}$$

and

$$\begin{aligned}\min_{m'\neq m}\Big[D_{KL}(f(\mathbf{x}^{(\mathbf{n})}|m,\tilde{\theta}_1,\tilde{\theta}_2)\|f(\mathbf{x}^{(\mathbf{n})}|m',\tilde{\theta}_1,\tilde{\theta}_2))\Big] &= \\ = \min_{m'\neq m}\{l\cdot D_{KL}(f_2(\cdot|\tilde{\theta}_2)\|f_1(\cdot|\tilde{\theta}_1)), l\cdot D_{KL}(f_1(\cdot|\tilde{\theta}_1)\|f_2(\cdot|\tilde{\theta}_2))\} & \\ = \min_{m'\neq m}\{D_{KL}(f_2(\cdot|\tilde{\theta}_2)\|f_1(\cdot|\tilde{\theta}_1)), D_{KL}(f_1(\cdot|\tilde{\theta}_1)\|f_2(\cdot|\tilde{\theta}_2))\}\cdot \underbrace{\min_{m'\neq m}\{l\}}_{1}.&\end{aligned} \tag{11}$$

We observe that equation (11) is only a function of $\tilde{\theta}_1$ and $\tilde{\theta}_2$ and does not depend on m. Thus, $\pi(m) \propto 1$ and, therefore,

$$\pi(m) = \frac{1}{n-1} \qquad m \in \{1,\ldots,n-1\}. \tag{12}$$

This prior was for example used in an econometric context by Koop and Potter (2009) with the rationale of giving equal weight to every possible change point location.

2.2 Multivariate Change Point Problem

In this section, we address the change point problem in its generality by assuming that there are $1 \leq k < n$ change points. In particular, for the data $\mathbf{x}^{(n)} = (x_1,\ldots,x_n)$, we consider the following sampling distribution

$$f(\mathbf{x}^{(n)}|\boldsymbol{m},\tilde{\boldsymbol{\theta}}) = \prod_{i=1}^{m_1} f_1(x_i|\tilde{\theta}_1)\prod_{j=1}^{k-1}\prod_{i=m_j+1}^{m_{j+1}} f_{j+1}(x_i|\tilde{\theta}_{j+1})\prod_{i=m_k+1}^{n} f_{k+1}(x_i|\tilde{\theta}_{k+1}), \tag{13}$$

where $\boldsymbol{m} = (m_1,\ldots,m_k)$, $1 \leq m_1 < m_2 < \ldots < m_k < n$, is the vector of the change point locations and $\tilde{\boldsymbol{\theta}} = (\tilde{\theta}_1,\ldots,\tilde{\theta}_k,\tilde{\theta}_{k+1})$ is the vector of the parameters of the underlying probability distributions. Schematically:

$$\begin{array}{llll} X_1 & ,\ldots, & X_{m_1}|\tilde{\theta}_1 & \overset{\text{i.i.d.}}{\sim} f_1(\cdot|\tilde{\theta}_1) \\ X_{m_1+1} & ,\ldots, & X_{m_2}|\tilde{\theta}_2 & \overset{\text{i.i.d.}}{\sim} f_2(\cdot|\tilde{\theta}_2) \\ \vdots & ,\ldots, & \vdots & \vdots\ldots\vdots \\ X_{m_{k-1}+1} & ,\ldots, & X_{m_k}|\tilde{\theta}_k & \overset{\text{i.i.d.}}{\sim} f_k(\cdot|\tilde{\theta}_k) \\ X_{m_k+1} & ,\ldots, & X_n|\tilde{\theta}_{k+1} & \overset{\text{i.i.d.}}{\sim} f_{k+1}(\cdot|\tilde{\theta}_{k+1}). \end{array}$$

If $f_1 = f_2 = \cdots = f_{k+1}$, then it is reasonable to assume $\tilde{\theta}_1 \neq \tilde{\theta}_2 \neq \cdots \neq \tilde{\theta}_k \neq \tilde{\theta}_{k+1}$. In a similar fashion to the single change point case, we cannot assume $m_k = n$ since we require exactly k change points.

In this case, due to the multivariate nature of the vector $\boldsymbol{m} = (m_1, \ldots, m_k)$, the derivation of the loss-based prior is not as straightforward as in the one dimensional case. In fact, the derivation of the prior is based on heuristic considerations supported by the below Theorem 1 (which proof is in the Appendix). In particular, we are able to prove an analogous of equations (9) and (10) when only one component is arbitrarily perturbed. Let us define the following functions:

$$d_j^{+1}(\tilde{\boldsymbol{\theta}}) = D_{KL}(f_{j+1}(\cdot|\tilde{\theta}_{j+1}) \| f_j(\cdot|\tilde{\theta}_j))$$
$$d_j^{-1}(\tilde{\boldsymbol{\theta}}) = D_{KL}(f_j(\cdot|\tilde{\theta}_j) \| f_{j+1}(\cdot|\tilde{\theta}_{j+1})),$$

where $j \in \{1, 2, \ldots, k\}$. The following Theorem is useful to understand the behaviour of the loss-based prior in the general case.

Theorem 1. *Let $f(\mathbf{x}^{(n)}|\boldsymbol{m}, \tilde{\boldsymbol{\theta}})$ be the sampling distribution defined in equation* (13) *and consider $j \in \{1, \ldots, k\}$. Let $\boldsymbol{m}'$ be such that $m'_i = m_i$ for $i \neq j$, and let the component m'_j be such that $m'_j \neq m_j$ and $m_{j-1} < m'_j < m_{j+1}$. Therefore,*

$$D_{KL}(f(\mathbf{x}^{(n)}|\boldsymbol{m}, \tilde{\boldsymbol{\theta}}) \| f(\mathbf{x}^{(n)}|\boldsymbol{m}', \tilde{\boldsymbol{\theta}}) = |m'_j - m_j| d_j^S(\tilde{\boldsymbol{\theta}}),$$

where $S = \operatorname{sgn}(m'_j - m_j)$.

Note that, Theorem 1 states that the minimum Kullback–Leibler divergence is achieved when $m'_j = m_j + 1$ or $m'_j = m_j - 1$. This result is not surprising since the Kullback–Leibler divergence measures the degree of similarity between two distributions. The smaller the perturbation caused by changes in one of the parameters is, the smaller the Kullback–Leibler divergence between the two distributions is. If now we consider the general case of having k change points, it is straightforward to see that the Kullback–Leibler divergence is minimized when only one of the components of the vector $\boldsymbol{m}$ is perturbed by (plus or minus) one unit. As such, the loss-based prior depends on the vector of parameters $\tilde{\boldsymbol{\theta}}$ only, as in the one-dimensional case, yielding the uniform prior for $\boldsymbol{m}$.

Therefore, the loss-based prior on the multivariate change point location is

$$\pi(\boldsymbol{m}) = \left\{ \binom{n-1}{k} \right\}^{-1}, \tag{14}$$

where $\boldsymbol{m} = (m_1, \ldots, m_k), 1 \leq m_1 < m_2 < \ldots < m_k < n$. The denominator in equation (14) has the above form because, for every number of k change points, we are interested in the number of k-subsets from a set of $n-1$ elements, which is $\binom{n-1}{k}$. The same prior was also derived in a different way by Girón et al. (2007).

3 Objective Prior on the Number of Change Points

Here, we approach the change point analysis as a model selection problem. In particular, we define an objective prior on the space of models, where each model represents a certain number of change points (including the case of no change points). The method adopted to define an objective prior on the space of models is the one introduced in Villa and Walker (2015*b*). We proceed as

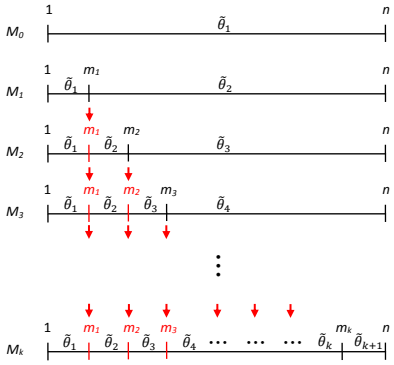

Figure 1: Diagram showing the way we specify our models. The arrows indicate that the respective change point locations remain fixed from the previous model to the current one.

follows. Assume we have to select from $k+1$ possible models. Let M_0 be the model with no change points, M_1 the model with one change point and so on. Generalising, model M_k corresponds to the model with k change points. The idea is that the current model encompasses the change point locations of the previous model. As an example, in model M_3 the first two change point locations will be the same as in the case of model M_2. To illustrate the way we envision our models, we have provided Figure 1. Keeping in line with the notation used during the introductory section, for the model M_k, with $k < n$ a non-negative integer, the model parameter θ_k is equivalent to the vector $(m_1, m_2, \ldots, m_k, \tilde{\theta}_1, \tilde{\theta}_2, \ldots, \tilde{\theta}_{k+1})$. Here, $\tilde{\theta}_1, \tilde{\theta}_2, \ldots, \tilde{\theta}_{k+1}$ represent the parameters of the underlying sampling distributions considered under model M_k, and $m_1, m_2, \ldots, m_k$ are the respective k change point locations, as in Figure 1.

Based on the way we have specified our models, which are in direct correspondence with the number of change points and their locations, we state Theorem 2 (which proof is in the Appendix).

Theorem 2. *Let*

$$D_{KL}(M_i \| M_j) = D_{KL}(f(\mathbf{x}^{(n)}|\theta_i) \| f(\mathbf{x}^{(n)}|\theta_j)).$$

For any $0 \leq i < j \leq k$ *integers, with* $k < n$*, and the convention* $m_{j+1} = n$*, we have the following:*

$$D_{KL}(M_i \| M_j) = \sum_{q=i+1}^{j} \left[(m_{q+1} - m_q) \cdot D_{KL}(f_{i+1}(\cdot|\tilde{\theta}_{i+1}) \| f_{q+1}(\cdot|\tilde{\theta}_{q+1})) \right],$$

and

$$D_{KL}(M_j \| M_i) = \sum_{q=i+1}^{j} \left[(m_{q+1} - m_q) \cdot D_{KL}(f_{q+1}(\cdot|\tilde{\theta}_{q+1}) \| f_{i+1}(\cdot|\tilde{\theta}_{i+1})) \right].$$

The result in Theorem 2 is useful when the model selection exercise is implemented. Indeed, the Villa and Walker (2015*b*) approach requires the computation of the Kullback–Leibler divergences in Theorem 2. Recalling equation (3), the objective priors on the model prior probabilities are given by:

$$\Pr(M_j) \propto \exp\left\{ \mathbb{E}_{\pi_j} \left[\inf_{\theta_i, i \neq j} D_{KL}(M_j \| M_i) \right] \right\} \qquad j = 0, 1, \ldots, k. \tag{15}$$

For illustrative purposes, in the Appendix we derive the model prior probabilities to perform model selection among M_0, M_1 and M_2.

Remark. In the case where the changes in the underlying sampling distribution are limited to the parameter values, the model prior probabilities defined in (15) follow the uniform distribution. That is, $\Pr(M_j) \propto 1$. In the real data example illustrated in Section 5.1, we indeed consider a problem where the above case occurs.

3.1 A special case: selection between M_0 and M_1

Let us consider the case where we have to estimate whether there is or not a change point in a set of observations. This implies that we have to choose between model M_0 (i.e. no change point) and M_1 (i.e. one change point). Following our approach, we have:

$$\Pr(M_0) \propto \exp\left\{\mathbb{E}_{\pi_0}\left[\inf_{\tilde{\theta}_2} D_{KL}(f_1(\cdot|\tilde{\theta}_1)\|f_2(\cdot|\tilde{\theta}_2))\right]\right\}, \tag{16}$$

and

$$\Pr(M_1) \propto \exp\left\{\mathbb{E}_{\pi_1}\left[(n-m_1)\cdot\inf_{\tilde{\theta}_1} D_{KL}(f_2(\cdot|\tilde{\theta}_2)\|f_1(\cdot|\tilde{\theta}_1))\right]\right\}. \tag{17}$$

Now, let us assume independence between the prior on the change point location and the prior on the parameters of the underlying sampling distributions, that is $\pi_1(m_1, \tilde{\theta}_1, \tilde{\theta}_2) = \pi_1(m_1)\pi_1(\tilde{\theta}_1, \tilde{\theta}_2)$. Let us further recall that, as per equation (14), $\pi_1(m_1) = 1/(n-1)$. As such, we observe that the model prior probability on M_1 becomes:

$$\Pr(M_1) \propto \exp\left\{\left(\frac{n}{2}\right)\mathbb{E}_{\pi_1(\tilde{\theta}_1,\tilde{\theta}_2)}\left[\inf_{\tilde{\theta}_1} D_{KL}(f_2(\cdot|\tilde{\theta}_2)\|f_1(\cdot|\tilde{\theta}_1))\right]\right\}. \tag{18}$$

We notice that the model prior probability for model M_1 is increasing when the sample size increases. This behaviour occurs whether there is or not a change point in the data.

We propose to address the above problem by using a non-uniform prior for m_1. A reasonable alternative, which works quite well in practice, would be to use the following shifted binomial as prior:

$$\pi_1(m_1) = \binom{n-2}{m_1-1}\left(\frac{n-1}{n}\right)^{m_1-1}\left(\frac{1}{n}\right)^{n-m_1-1}, 1 \leq m_1 \leq n-1. \tag{19}$$

To argument the choice of (19), we note that as n increases, the probability mass will be more and more concentrated towards the upper end of the support. Therefore, from equations (17) and (19) follows:

$$\Pr(M_1) \propto \exp\left\{\left(\frac{2n-2}{n}\right)\mathbb{E}_{\pi_1(\tilde{\theta}_1,\tilde{\theta}_2)}\left[\inf_{\tilde{\theta}_1} D_{KL}(f_2(\cdot|\tilde{\theta}_2)\|f_1(\cdot|\tilde{\theta}_1))\right]\right\}. \tag{20}$$

For the more general case where one considers more than two models, the problem highlighted in equation (18) vanishes.

4 Change Point Analysis on Simulated Data

In this section, we present the results of several simulation studies based on the methodologies discussed in Sections 2 and 3. We start with a scenario involving discrete distributions in the context of the one change point

problem. We then show the results obtained when we consider continuous distributions for the case of two change points. The choice of the underlying sampling distribution is in line with Villa and Walker (2015*b*).

Scenario 1. The first scenario concerns the choice between models M_0 and M_1. Specifically, for M_0 we have:

$$X_1, X_2, \ldots, X_n | p \overset{\text{i.i.d.}}{\sim} \text{Geometric}(p)$$

and for M_1 we have:

$$X_1, X_2, \ldots, X_{m_1} | p \overset{\text{i.i.d.}}{\sim} \text{Geometric}(p)$$
$$X_{m_1+1}, X_{m_1+2}, \ldots, X_n | \lambda \overset{\text{i.i.d.}}{\sim} \text{Poisson}(\lambda)$$

Let us denote with $f_1(\cdot|p)$ and $f_2(\cdot|\lambda)$, the probability mass functions of the Geometric and the Poisson distributions, respectively. The priors for the parameters of f_1 and f_2 are $p \sim \text{Beta}(a, b)$ and $\lambda \sim \text{Gamma}(c, d)$.

In the first simulation, we sample $n = 100$ observations from model M_0 with $p = 0.8$. To perform the change point analysis, we have chosen the following parameters for the priors on p and λ: $a = 2$, $b = 2$, $c = 3$ and $d = 1$. Applying the approach introduced in Section 3, we obtain $\Pr(M_0) \propto 1.59$ and $\Pr(M_1) \propto 1.81$. These model priors yield the posterior distribution probabilities (refer to equation (4)) $\Pr(M_0|\mathbf{x}^{(n)}) = 0.92$ and $\Pr(M_1|\mathbf{x}^{(n)}) = 0.08$. As expected, the selection process strongly indicates the true model as M_0. Table 1 reports the above probabilities including other information, such as the appropriate Bayes factors.

The second simulation looked at the opposite set up, that is we sample $n = 100$ observations from M_1, with $p = 0.8$ and $\lambda = 3$. We have sampled 50 data points from the Geometric distribution and the remaining 50 data points from the Poisson distribution. In Figure 2 we have plotted the simulated sample, where it is legitimate to assume a change in the underlying distribution. Using the same prior parameters as above, we obtain $\Pr(M_0|\mathbf{x}^{(n)}) = 0.06$ and $\Pr(M_1|\mathbf{x}^{(n)}) = 0.94$. Again, the model selection process is assigning heavy posterior mass to the true model M_1. These results are further detailed in Table 1.

We have also performed the same analyis where we have set M_0 as a Poisson and M_1 as a combination of a Poisson as first distribution and of a Geometric as the second distribution. The simulation results are not reported here, however we have not detected any noticeable difference from the previous setup.

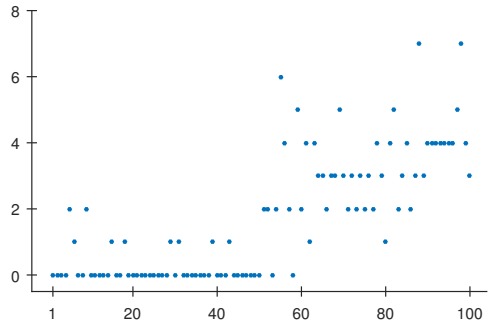

Figure 2: Scatter plot of the data simulated from model M_1 in Scenario 1.

	True model	
	M_0	M_1
$\Pr(M_0)$	0.47	0.47
$\Pr(M_1)$	0.53	0.53
B_{01}	12.39	0.08
B_{10}	0.08	12.80
$\Pr(M_0\|\mathbf{x}^{(n)})$	0.92	0.06
$\Pr(M_1\|\mathbf{x}^{(n)})$	0.08	0.94

Table 1: Model prior, Bayes factor and model posterior probabilities for the change point analysis in Scenario 1. We considered samples from, respectively, model M_0 and model M_1.

Scenario 2. In this scenario we consider the case where we have to select among three models, that is model M_0:

$$X_1, X_2, \ldots, X_n|\lambda, \kappa \overset{\text{i.i.d.}}{\sim} \text{Weibull}(\lambda, \kappa), \tag{21}$$

model M_1:

$$\begin{aligned} X_1, X_2, \ldots, X_{m_1}|\lambda, \kappa &\overset{\text{i.i.d.}}{\sim} \text{Weibull}(\lambda, \kappa) \\ X_{m_1+1}, X_{m_1+2}, \ldots, X_n|\mu, \tau &\overset{\text{i.i.d.}}{\sim} \text{Log-normal}(\mu, \tau), \end{aligned} \tag{22}$$

with $1 \leq m_1 \leq n-1$ being the location of the single change point, and model M_2:

$$\begin{aligned} X_1, X_2, \ldots, X_{m_1}|\lambda, \kappa &\overset{\text{i.i.d.}}{\sim} \text{Weibull}(\lambda, \kappa) \\ X_{m_1+1}, X_{m_1+2}, \ldots, X_{m_2}|\mu, \tau &\overset{\text{i.i.d.}}{\sim} \text{Log-normal}(\mu, \tau) \\ X_{m_2+1}, X_{m_2+2}, \ldots, X_n|\alpha, \beta &\overset{\text{i.i.d.}}{\sim} \text{Gamma}(\alpha, \beta), \end{aligned} \tag{23}$$

with $1 \leq m_1 < m_2 \leq n-1$ representing the locations of the two change points, such that m_1 corresponds exactly to the same location as in model M_1. Analogously to the previous scenario, we sample from each model in turn and perform the selection to detect the number of change points.

Let $f_1(\cdot|\lambda,\kappa)$, $f_2(\cdot|\mu,\tau)$ and $f_3(\cdot|\alpha,\beta)$ represent the Weibull, Log-normal and Gamma densities, respectively, with $\tilde{\theta}_1 = (\lambda,\kappa)$, $\tilde{\theta}_2 = (\mu,\tau)$ and $\tilde{\theta}_3 = (\alpha,\beta)$. We assume a Normal prior on μ and Gamma priors on all the other parameters as follows:

$$\lambda \sim \text{Gamma}(1.5,1) \qquad \kappa \sim \text{Gamma}(5,1) \qquad \mu \sim \text{Normal}(0.05,1),$$

$$\tau \sim \text{Gamma}(16,1) \qquad \alpha \sim \text{Gamma}(10,1) \qquad \beta \sim \text{Gamma}(0.2,0.1).$$

In the first exercise, we have simulated $n = 100$ observations from model M_0, where we have set $\lambda = 1.5$ and $\kappa = 5$. The model priors we obtain are $\Pr(M_0) \propto 1.09$, $\Pr(M_1) \propto 1.60$ and $\Pr(M_2) \propto 1.37$, yielding the posteriors $\Pr(M_0|\mathbf{x}^{(n)}) = 0.96$, $\Pr(M_1|\mathbf{x}^{(n)}) = 0.04$ and $\Pr(M_2|\mathbf{x}^{(n)}) = 0.00$. We then see that the approach assigns high mass to the true model M_0. Table 2 reports the above probabilities and the corresponding Bayes factors. The

	True model		
	M_0	M_1	M_2
$\Pr(M_0)$	0.27	0.27	0.27
$\Pr(M_1)$	0.39	0.39	0.39
$\Pr(M_2)$	0.34	0.34	0.34
B_{01}	36.55	3.24×10^{-4}	4.65×10^{-40}
B_{02}	1.84^3	0.02	1.27×10^{-45}
B_{12}	50.44	55	2.72×10^{-6}
$\Pr(M_0\|\mathbf{x}^{(n)})$	0.96	0.00	0.00
$\Pr(M_1\|\mathbf{x}^{(n)})$	0.04	0.98	0.00
$\Pr(M_2\|\mathbf{x}^{(n)})$	0.00	0.02	1.00

Table 2: Model prior, Bayes factor and model posterior probabilities for the change point analysis in Scenario 2. We considered samples from, respectively, model M_0, model M_1 and model M_2.

second simulation was performed by sampling 50 observations from a Weibull with parameter values as in the previous exercise, and the remaining 50 observations from a Log-normal density with location parameter $\mu = 0.05$ and scale parameter $\tau = 16$. The data is displayed in Figure 3.

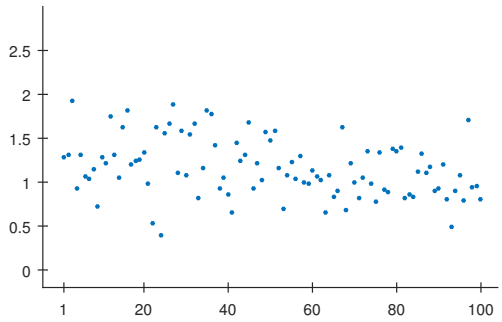

Figure 3: Scatter plot of the observations simulated from model M_1 in Scenario 2.

The model posterior probabilities are $\Pr(M_0|\mathbf{x}^{(n)}) = 0.00$, $\Pr(M_1|\mathbf{x}^{(n)}) = 0.98$ and $\Pr(M_2|\mathbf{x}^{(n)}) = 0.02$, which are reported in Table 2. In this case as well, we see that the model selection procedure indicates M_1 as the true model, as expected.

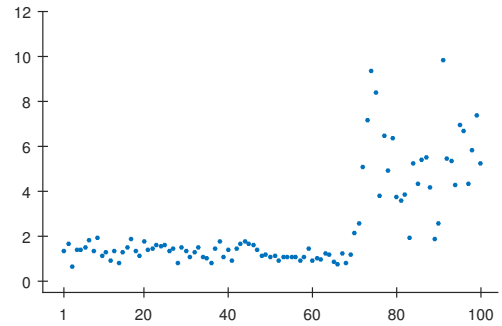

Figure 4: Scatter plot of the observations simulated from model M_2 in Scenario 2.

Finally, for the third simulation exercise we sample 50 and 20 data points from, respectively, a Weibull and a Log-normal with parameter values as defined above, and the last 30 observations are sampled from a Gamma distribution with parameters $\alpha = 10$ and $\beta = 2$. From Table 2 we note that the posterior distribution on the model space accumulates on the true model M_2.

5 Change Point Analysis on Real Data

In this section we illustrate the proposed approach applied to real data. We first consider a well known data set which has been extensively studied in

the literature of the change point analysis, that is the British coal-mining disaster data (Carlin et al., 1992). The second set of data we consider refers to the daily returns of the S&P 500 index observed over a period of four years. The former data set will be investigated in Section 5.1, while the latter in Section 5.2.

5.1 British Coal-Mining Disaster Data

The British coal-mining disaster data consists of the yearly number of deaths for the British coal miners over the period 1851-1962. It is believed that the change in the working conditions, and in particular in the enhancement of the security measures, led to a decrease in the number of deaths. This calls for a model which can take into account a change in the underlying distribution around a certain observed year. With the proposed methodology we wish to detect if the assumption is appropriate. In particular, if a model with one change point is more suitable to represent the data than a model where no changes in the sampling distribution are assumed. Figure 5 shows the number of deaths per year in the British coal-mining industry from 1851 to 1962. As in Chib (1998), we assume a Poisson sampling distribution with a

Figure 5: Scatter plot of the British coal-mining disaster data.

possible change in the parameter value. That is

$$\begin{aligned} X_1, X_2, \ldots, X_m | \phi_1 &\overset{\text{i.i.d.}}{\sim} \text{Poisson}(\phi_1) \\ X_{m+1}, X_{m+2}, \ldots, X_n | \phi_2 &\overset{\text{i.i.d.}}{\sim} \text{Poisson}(\phi_2), \end{aligned} \tag{24}$$

where m is the unknown location of the single change point, such that $1 \leq m \leq n$, and a Gamma$(2, 1)$ is assumed for ϕ_1 and ϕ_2. The case $m = n$ corresponds to the scenario with no change point, that is model M_0. The case $m < n$ assumes one change point, that is model M_1.

Let $f_1(\cdot|\phi_1)$ and $f_2(\cdot|\phi_2)$ be the Poisson distributions with parameters ϕ_1 and ϕ_2, respectively. Then, the analysis is performed by selecting between model M_0, that is when the sampling distribution is f_1, and model M_1, where the sampling distribution is f_1 up to a certain $m < n$ and f_2 from $m+1$ to n.

As highlighted in the Remark at the end of Section 3, the prior on the model space is the uniform, that is $\Pr(M_0) = \Pr(M_1) = 0.5$. The proposed model selection approach, leads to the Bayes factors $B_{01} = 1.61 \times 10^{-13}$ and $B_{10} = 6.20 \times 10^{12}$, where it is obvious that the odds are strongly in favour of model M_1. Indeed, we have $\Pr(M_1|\mathbf{x}^{(n)}) \approx 1$.

5.2 Daily S&P 500 Absolute Log-Return Data

The second real data analysis aims to detect change points in the absolute value of the daily logarithmic returns of the S&P500 index observed from the 14/01/2008 to the 31/12/2011 (see Figure 6). As underlying sampling distributions we consider the Weibull and the Log-normal (Yu, 2001), and the models among which we select are as follows. M_0 is a Weibull(λ, κ), M_1 is formed by a Weibull(λ, κ) and a Log-normal(μ_1, τ_1) and, finally, M_2 is formed by a Weibull(λ, κ), a Log-normal(μ_1, τ_1) and a Log-normal(μ_2, τ_2). An interesting particularity of this problem is that we will consider a scenario where the changes are in the underlying distribution or in the parameter values of the same distribution. Due to the large sample size of the data set, we have approximated the Bayes factor using the Schwarz criterion (Kass and Raftery, 1995). Therefore, in this case the specification of the priors for the parameters of the underlying distributions is not necessary. From the results

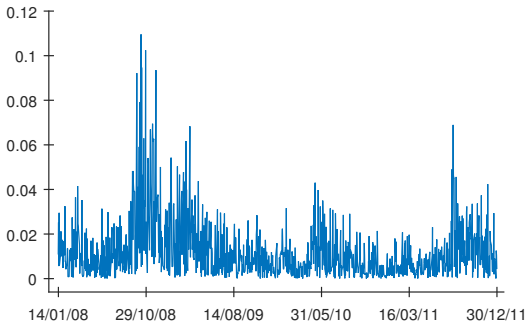

Figure 6: Absolute daily log-returns of the S&P500 index from 14/01/08 to 30/12/11.

in Table 3, we see that the model indicated by the proposed approach is M_2. In other words, there is very strong indication that there are two change

points in the data set. From Table 3 we note that the prior on model M_1 and

$\Pr(M_0)$	0.36
$\Pr(M_1)$	0.32
$\Pr(M_2)$	0.32
B_{01}	7.72×10^{18}
B_{02}	3.30×10^{-3}
B_{12}	4.28×10^{-22}
$\Pr(M_0\|\mathbf{x}^{(n)})$	0.00
$\Pr(M_1\|\mathbf{x}^{(n)})$	0.00
$\Pr(M_2\|\mathbf{x}^{(n)})$	1.00

Table 3: Model prior, Bayes factor and model posterior probabilities for the S&P500 change point analysis.

M_2 assigned by the proposed method, are the same. This is not surprising as the only difference between the two models is an additional Log-normal distribution with different parameter values.

6 Conclusion

Bayesian inference in change point problems under the assumption of minimal prior information has not been deeply explored in the past, as the limited literature on the matter shows.

We contribute to the area by deriving an objective prior distribution to detect change point locations, when the number of change points is known a priori. As a change point location can be interpreted as a discrete parameter, we apply recent results in the literature (Villa and Walker, 2015*a*) to make inference. The resulting prior distribution, which is uniform, it is not new in the literature (Girón et al., 2007), and therefore can be considered as a validation of the proposed approach.

A second major contribution is in defining an objective prior on the number of change points, which has been approached by considering the problem as a model selection exercise. The results of the proposed method on both simulated and real data, allow to show the strength of the approach in estimating the number of change points in a series of observations. A point to note is in the generality of the scenario considered. Indeed, we consider situations where the change is in the value of the parameter(s) of the underlying sampling distribution, or in the distribution itself. Of particular interest, is the

last real data analysis (S&P 500 index), where we consider a scenario where we have both types of changes, that is the distribution for the first change point and on the parameters of the distribution for the second.

References

Berk, R. H. (1966), 'Limiting Behavior of Posterior Distributions when the Model is Incorrect', *The Annals of Mathematical Statistics* **37**(1), 51–58.

Carlin, B. P., Gelfand, A. E. and Smith, A. F. (1992), 'Hierarchical Bayesian Analysis of Changepoint Problems', *Applied Statistics* **41**(2), 389–405.

Chib, S. (1998), 'Estimation and Comparison of Multiple Change-point Models', *Journal of Econometrics* **86**(2), 221 – 241.

Girón, F. J., Moreno, E. and Casella, G. (2007), Objective Bayesian Analysis of Multiple Changepoints for Linear Models, *in* J. Bernardo, M. Bayarri, J. Berger, A. Dawid, D. Heckerman, A. Smith and M. West, eds, 'Bayesian Statistics 8', Oxford University Press, pp. 227–252.

Kass, R. E. and Raftery, A. E. (1995), 'Bayes Factors', *Journal of the American Statistical Association* **90**(430), 773–795.

Koop, G. and Potter, S. M. (2009), 'Prior Elicitation in Multiple Change-Point Models', *International Economic Review* **50**(3), 751–772.

Kullback, S. and Leibler, R. A. (1951), 'On Information and Sufficiency', *The Annals of Mathematical Statistics* **22**(1), 79–86.

Moreno, E., Casella, G. and Garcia-Ferrer, A. (2005), 'An Objective Bayesian Analysis of the Change Point Problem', *Stochastic Environmental Research and Risk Assessment* **19**(3), 191–204.

Villa, C. and Walker, S. (2015*a*), 'An Objective Approach to Prior Mass Functions for Discrete Parameter Spaces', *Journal of the American Statistical Association* **110**(511), 1072–1082.

Villa, C. and Walker, S. (2015*b*), 'An Objective Bayesian Criterion to Determine Model Prior Probabilities', *Scandinavian Journal of Statistics* **42**(4), 947–966.

Yu, J. (2001), Chapter 6 - Testing for a Finite Variance in Stock Return Distributions, *in* J. Knight, and S. Satchell, eds, 'Return Distributions in Finance', Quantitative Finance, Butterworth-Heinemann, Oxford, pp. 143 – 164.

Appendix

A Model prior probabilities to select among models M_0, M_1 and M_2

Here we show how model prior probabilities can be derived for the relatively simple case of selecting among scenarios with no change points (M_0), one change point (M_1) or two change points (M_2). First, by applying the result in Theorem 2, we derive the Kullback–Leibler divergences between any two models. That is:

- the prior probability for model M_0 depends on the following quantities:

$$\begin{aligned} D_{KL}(M_0\|M_1) =& (n-m_1)\cdot D_{KL}(f_1(\cdot|\tilde{\theta}_1)\|f_2(\cdot|\tilde{\theta}_2)) \\ D_{KL}(M_0\|M_2) =& (m_2-m_1)\cdot D_{KL}(f_1(\cdot|\tilde{\theta}_1)\|f_2(\cdot|\tilde{\theta}_2)) \\ &+ (n-m_2)\cdot D_{KL}(f_1(\cdot|\tilde{\theta}_1)\|f_3(\cdot|\tilde{\theta}_3)) \end{aligned}$$

- the prior probability for model M_1 depends on the following quantities:

$$\begin{aligned} D_{KL}(M_1\|M_2) =& (n-m_2)\cdot D_{KL}(f_2(\cdot|\tilde{\theta}_2)\|f_3(\cdot|\tilde{\theta}_3)) \\ D_{KL}(M_1\|M_0) =& (n-m_1)\cdot D_{KL}(f_2(\cdot|\tilde{\theta}_2)\|f_1(\cdot|\tilde{\theta}_1)) \end{aligned}$$

- the prior probability for model M_2 depends on the following quantities:

$$\begin{aligned} D_{KL}(M_2\|M_1) =& (n-m_2)\cdot D_{KL}(f_3(\cdot|\tilde{\theta}_3)\|f_2(\cdot|\tilde{\theta}_2)) \\ D_{KL}(M_2\|M_0) =& (m_2-m_1)\cdot D_{KL}(f_2(\cdot|\tilde{\theta}_2)\|f_1(\cdot|\tilde{\theta}_1)) \\ &+ (n-m_2)\cdot D_{KL}(f_3(\cdot|\tilde{\theta}_3)\|f_1(\cdot|\tilde{\theta}_1)) \end{aligned}$$

The next step is to derive the minimum Kullback–Leibler divergence computed at each model:

- for model M_0:

$$\begin{aligned}
\inf_{\theta_1} D_{KL}(M_0\|M_1) &= \underbrace{\left[\inf_{m_1\neq n}(n-m_1)\right]}_{1}\cdot\left[\inf_{\tilde{\theta}_2} D_{KL}(f_1(\cdot|\tilde{\theta}_1)\|f_2(\cdot|\tilde{\theta}_2))\right] \\
&= \inf_{\tilde{\theta}_2} D_{KL}(f_1(\cdot|\tilde{\theta}_1)\|f_2(\cdot|\tilde{\theta}_2)) \\
\inf_{\theta_2} D_{KL}(M_0\|M_2) &= \underbrace{\left[\inf_{m_1\neq m_2}(m_2-m_1)\right]}_{1}\cdot\left[\inf_{\tilde{\theta}_2} D_{KL}(f_1(\cdot|\tilde{\theta}_1)\|f_2(\cdot|\tilde{\theta}_2))\right] \\
&\quad + \underbrace{\left[\inf_{m_2\neq n}(n-m_2)\right]}_{1}\cdot\left[\inf_{\tilde{\theta}_3} D_{KL}(f_1(\cdot|\tilde{\theta}_1)\|f_3(\cdot|\tilde{\theta}_3))\right] \\
&= \inf_{\tilde{\theta}_2} D_{KL}(f_1(\cdot|\tilde{\theta}_1)\|f_2(\cdot|\tilde{\theta}_2)) + \inf_{\tilde{\theta}_3} D_{KL}(f_1(\cdot|\tilde{\theta}_1)\|f_3(\cdot|\tilde{\theta}_3))
\end{aligned}$$

- for model M_1:

$$\begin{aligned}
\inf_{\theta_2} D_{KL}(M_1\|M_2) &= \underbrace{\left[\inf_{m_2\neq n}(n-m_2)\right]}_{1}\cdot\left[\inf_{\tilde{\theta}_3} D_{KL}(f_2(\cdot|\tilde{\theta}_2)\|f_3(\cdot|\tilde{\theta}_3))\right] \\
&= \inf_{\tilde{\theta}_3} D_{KL}(f_2(\cdot|\tilde{\theta}_2)\|f_3(\cdot|\tilde{\theta}_3)) \\
\inf_{\theta_0=\tilde{\theta}_1} D_{KL}(M_1\|M_0) &= (n-m_1)\cdot\inf_{\tilde{\theta}_1} D_{KL}(f_2(\cdot|\tilde{\theta}_2)\|f_1(\cdot|\tilde{\theta}_1))
\end{aligned}$$

- for model M_2:

$$\begin{aligned}
\inf_{\theta_1} D_{KL}(M_2\|M_1) &= (n-m_2)\cdot\inf_{\tilde{\theta}_2} D_{KL}(f_3(\cdot|\tilde{\theta}_3)\|f_2(\cdot|\tilde{\theta}_2)) \\
\inf_{\theta_0=\tilde{\theta}_1} D_{KL}(M_2\|M_0) &= (m_2-m_1)\cdot\inf_{\tilde{\theta}_1} D_{KL}(f_2(\cdot|\tilde{\theta}_2)\|f_1(\cdot|\tilde{\theta}_1)) \\
&\quad + (n-m_2)\cdot\inf_{\tilde{\theta}_1} D_{KL}(f_3(\cdot|\tilde{\theta}_3)\|f_1(\cdot|\tilde{\theta}_1))
\end{aligned}$$

Therefore, the model prior probabilities can be computed through equation (15), so that:

- the model prior probability $\Pr(M_0)$ is proportional to the exponential of the minimum between:

$$\left\{\mathbb{E}_{\pi_0}\left[\inf_{\tilde{\theta}_2} D_{KL}(f_1(\cdot|\tilde{\theta}_1)\|f_2(\cdot|\tilde{\theta}_2))\right], \mathbb{E}_{\pi_0}\left[\inf_{\tilde{\theta}_2} D_{KL}(f_1(\cdot|\tilde{\theta}_1)\|f_2(\cdot|\tilde{\theta}_2)) + \inf_{\tilde{\theta}_3} D_{KL}(f_1(\cdot|\tilde{\theta}_1)\|f_3(\cdot|\tilde{\theta}_3))\right]\right\}$$

- the model prior probability $\Pr(M_1)$ is proportional to the exponential of the minimum between:

$$\left\{ \mathbb{E}_{\pi_1} \left[\inf_{\tilde{\theta}_3} D_{KL}(f_2(\cdot|\tilde{\theta}_2) \| f_3(\cdot|\tilde{\theta}_3)) \right], \right.$$
$$\left. \mathbb{E}_{\pi_1} \left[(n - m_1) \cdot \inf_{\tilde{\theta}_1} D_{KL}(f_2(\cdot|\tilde{\theta}_2) \| f_1(\cdot|\tilde{\theta}_1)) \right] \right\}$$

- the model prior probability $\Pr(M_2)$ is proportional to the exponential of the minimum between:

$$\left\{ \mathbb{E}_{\pi_2} \left[(n - m_2) \cdot \inf_{\tilde{\theta}_2} D_{KL}(f_3(\cdot|\tilde{\theta}_3) \| f_2(\cdot|\tilde{\theta}_2)) \right], \right.$$
$$\mathbb{E}_{\pi_2} \left[(m_2 - m_1) \cdot \inf_{\tilde{\theta}_1} D_{KL}(f_2(\cdot|\tilde{\theta}_2) \| f_1(\cdot|\tilde{\theta}_1)) + (n - m_2) \right.$$
$$\left. \left. \cdot \inf_{\tilde{\theta}_1} D_{KL}(f_3(\cdot|\tilde{\theta}_3) \| f_1(\cdot|\tilde{\theta}_1)) \right] \right\}$$

B Proofs

Proof of Theorem 1

We distinguish two cases: $S = +1$ and $S = -1$. When $S = +1$, equivalent to $m_j < m'_j$:

$$D_{KL}(f(\mathbf{x}^{(n)}|\boldsymbol{m},\tilde{\boldsymbol{\theta}})\|f(\mathbf{x}^{(n)}|\boldsymbol{m}',\tilde{\boldsymbol{\theta}})) = \int f(\mathbf{x}^{(n)}|\boldsymbol{m},\tilde{\boldsymbol{\theta}}) \cdot \ln\left(\frac{f(\mathbf{x}^{(n)}|\boldsymbol{m},\tilde{\boldsymbol{\theta}})}{f(\mathbf{x}^{(n)}|\boldsymbol{m}',\tilde{\boldsymbol{\theta}})}\right) \mathrm{d}\mathbf{x}^{(n)}$$

$$= \int f(\mathbf{x}^{(n)}|\boldsymbol{m},\tilde{\boldsymbol{\theta}}) \cdot \left[\sum_{i=m_j+1}^{m'_j} \ln\left(\frac{f_{j+1}(x_i|\tilde{\theta}_{j+1})}{f_j(x_i|\tilde{\theta}_j)}\right)\right] \mathrm{d}\mathbf{x}^{(n)}$$

$$= \sum_{i=m_j+1}^{m'_j} \int f(\mathbf{x}^{(n)}|\boldsymbol{m},\tilde{\boldsymbol{\theta}}) \cdot \left[\ln\left(\frac{f_{j+1}(x_i|\tilde{\theta}_{j+1})}{f_j(x_i|\tilde{\theta}_j)}\right)\right] \mathrm{d}\mathbf{x}^{(n)}$$

$$= \sum_{i=m_j+1}^{m'_j} \left\{1^{n-1} \cdot \int f_{j+1}(x_i|\tilde{\theta}_{j+1}) \cdot \left[\ln\left(\frac{f_{j+1}(x_i|\tilde{\theta}_{j+1})}{f_j(x_i|\tilde{\theta}_j)}\right)\right] \mathrm{d}x_i\right\}$$

$$= \sum_{i=m_j+1}^{m'_j} D_{KL}(f_{j+1}(x_i|\tilde{\theta}_{j+1})\|f_j(x_i|\tilde{\theta}_j))$$

$$=(m'_j - m_j) \cdot D_{KL}(f_{j+1}(\cdot|\tilde{\theta}_{j+1})\|f_j(\cdot|\tilde{\theta}_j))$$

$$=(m'_j - m_j) \cdot d_j^{+1}(\tilde{\boldsymbol{\theta}}). \tag{25}$$

When $S = -1$, equivalent to $m_j > m'_j$, in a similar fashion, we get

$$D_{KL}(f(\mathbf{x}^{(n)}|\boldsymbol{m},\tilde{\boldsymbol{\theta}})\|f(\mathbf{x}^{(n)}|\boldsymbol{m}',\tilde{\boldsymbol{\theta}})) = (m_j - m'_j) \cdot d_j^{-1}(\tilde{\boldsymbol{\theta}}) \tag{26}$$

From equations (25) and (26), we get the result in Theorem 1.

Proof of Theorem 2

We recall that the model parameter θ_i is the vector $(m_1, m_2, \ldots, m_i, \tilde{\theta}_1, \tilde{\theta}_2, \ldots, \tilde{\theta}_{i+1})$, where $i = 0, 1, \ldots, k$. Here, $\tilde{\theta}_1, \tilde{\theta}_2, \ldots, \tilde{\theta}_{i+1}$ represent the parameters of the underlying sampling distributions considered under model M_i and $m_1, m_2, \ldots, m_i$ are the respective i change point locations. In this setting,

$$f(\mathbf{x}^{(n)}|\theta_i) = \prod_{r=1}^{m_1} f_1(x_r|\tilde{\theta}_1) \prod_{t=1}^{i-1} \prod_{r=m_t+1}^{m_{t+1}} f_{t+1}(x_r|\tilde{\theta}_{t+1}) \prod_{r=m_i+1}^{n} f_{i+1}(x_r|\tilde{\theta}_{i+1}) \tag{27}$$

We proceed to the computation of $D_{KL}(M_i \| M_j)$, that is the Kullback–Leibler divergence introduced in Section 3. Similarly to the proof of Theorem 1, we obtain the following result.

$$\begin{aligned} D_{KL}(M_i \| M_j) = & \sum_{r=m_{i+1}+1}^{m_{i+2}} \int f(\mathbf{x}^{(n)}|\theta_i) \ln \left(\frac{f_{i+1}(x_r|\tilde{\theta}_{i+1})}{f_{i+2}(x_r|\tilde{\theta}_{i+2})} \right) \mathrm{d}\mathbf{x}^{(n)} \\ & + \sum_{r=m_{i+2}+1}^{m_{i+3}} \int f(\mathbf{x}^{(n)}|\theta_i) \ln \left(\frac{f_{i+1}(x_r|\tilde{\theta}_{i+1})}{f_{i+3}(x_r|\tilde{\theta}_{i+3})} \right) \mathrm{d}\mathbf{x}^{(n)} + \\ & \ldots + \sum_{r=m_j+1}^{n} \int f(\mathbf{x}^{(n)}|\theta_i) \ln \left(\frac{f_{i+1}(x_r|\tilde{\theta}_{i+1})}{f_{j+1}(x_r|\tilde{\theta}_{j+1})} \right) \mathrm{d}\mathbf{x}^{(n)}. \end{aligned}$$

Given equation (27), if we integrate out the variables not involved in the logarithms, we obtain

$$\begin{aligned} D_{KL}(M_i \| M_j) = & (m_{i+2} - m_{i+1}) \cdot D_{KL}(f_{i+1}(\cdot|\tilde{\theta}_{i+1}) \| f_{i+2}(\cdot|\tilde{\theta}_{i+2})) \\ & + (m_{i+3} - m_{i+2}) \cdot D_{KL}(f_{i+1}(\cdot|\tilde{\theta}_{i+1}) \| f_{i+3}(\cdot|\tilde{\theta}_{i+3})) + \\ & \ldots + (n - m_j) \cdot D_{KL}(f_{i+1}(\cdot|\tilde{\theta}_{i+1}) \| f_{j+1}(\cdot|\tilde{\theta}_{j+1})). \end{aligned}$$

In a similar fashion, it can be shown that

$$\begin{aligned} D_{KL}(M_j \| M_i) = & (m_{i+2} - m_{i+1}) \cdot D_{KL}(f_{i+2}(\cdot|\tilde{\theta}_{i+2}) \| f_{i+1}(\cdot|\tilde{\theta}_{i+1})) \\ & + (m_{i+3} - m_{i+2}) \cdot D_{KL}(f_{i+3}(\cdot|\tilde{\theta}_{i+3}) \| f_{i+1}(\cdot|\tilde{\theta}_{i+1})) + \\ & \ldots + (n - m_j) \cdot D_{KL}(f_{j+1}(\cdot|\tilde{\theta}_{j+1}) \| f_{i+1}(\cdot|\tilde{\theta}_{i+1})) \end{aligned}$$